

NISTIR 88-4019



On the Analysis of Computer Performance Data

Jack C. M. Wang
John M. Gary*
Hari K. Iyer

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
(Formerly National Bureau of Standards)
Center for Computing and Applied Mathematics
Statistical Engineering Division
*Scientific Computing Division
Boulder, Colorado 80303-3328

January 1989

NISTIR 88-4019

On the Analysis of Computer Performance Data

Jack C. M. Wang
John M. Gary*
Hari K. Iyer

U.S. DEPARTMENT OF COMMERCE
National Institute of Standards and Technology
(Formerly National Bureau of Standards)
Center for Computing and Applied Mathematics
Statistical Engineering Division
*Scientific Computing Division
Boulder, Colorado 80303-3328

January 1989



National Bureau of Standards became the National Institute of Standards and Technology on August 23, 1988, when the Omnibus Trade and Competitiveness Act was signed. NIST retains all NBS functions. Its new programs will encourage improved use of technology by U.S. industry.

U.S. DEPARTMENT OF COMMERCE
Robert A. Mosbacher, Secretary
Ernest Ambler, Acting Under Secretary
for Technology
NATIONAL INSTITUTE OF STANDARDS
AND TECHNOLOGY
Raymond G. Kammer, Acting Director

On the Analysis of Computer Performance Data ¹

Jack C. M. Wang and John M. Gary
National Institute of Standards and Technology
Boulder, CO 80303

Hari K. Iyer
Department of Statistics
Colorado State University
Fort Collins, CO 80523
and
National Institute of Standards and Technology

This paper is devoted to an analysis of the data from the Livermore loops benchmark. We will show that in a general predictive sense the dimension of this data is rather small; perhaps between two and five. Two techniques are used to reduce the 72 loops timings for each machine to a few scores which characterize the machine. The first is based on a principal component analysis, the second on a cluster analysis of the loops. The validity of the reduction of the data to a lesser dimension is checked by various methods.

Key words: benchmarks; computers; performance; Livermore loops; principal components; clusters.

¹Any mention of a specific commercial product in this paper does not imply an endorsement by the National Institute of Standards and Technology

1 Introduction

This paper is concerned with the analysis of the benchmark data obtained from the Livermore Fortran Kernels (LFK). Our objective is to summarize this data and present it in a simple, clear manner with minimal loss of information. A related objective is to estimate the information content of the LFK loop timings.

The LFK code consists of 24 short Fortran code segments along with a driver to execute and time the segments and present the data in a standard format (McMahon[1986]). The timing results are given in megaflops (million floating point operations per second). The segments are all compute bound, no attempt is made to measure I/O rates. The segments consist primarily of short DO loops which are designed to cover a very wide range of execution rates. Therefore, the set contains a loop which will vectorize very well and may run at 1000 megaflops on a certain system, as well as another loop which may run at 4 megaflops on the same system. The loops in the 24 segments are each run with three different lengths, thus giving a total of 72 different test segments. We have selected 48 different machine/compiler systems from the data given in the report by McMahon and used this set to test our data analysis techniques. There is extensive experience with these segments on serial and vector machines, but very few results have been reported for parallel machines. We will refer to the 24 code segments as the "loops", and the 72 numbers obtained from timing the loops as the "loop runs".

The LFK benchmark data set describes each system by 72 numbers. Our objective is to describe each system by far fewer numbers, perhaps two to four. We will refer to these numbers as the "scores" for each system. McMahon's report reduces this data to two scores, the harmonic and geometric means of the megaflops rates. In addition, he sometimes adds the arithmetic mean to give three numbers to characterize the systems. It is certainly desirable that these numbers have an easily understood meaning; for example, one number might be a mean megaflops rate for loops which vectorize well, and another the rate for loops which do not vectorize. One way to do this is to divide the loops into two or three

groups and characterize the systems by the geometric mean megaflops rates over these groups. One group might contain the “fast” loops which vectorize easily, another those which vectorize poorly or not at all. Another method is to include all the loops in each group, but weight the loops differently in each group. McMahon’s paper gives megaflops rates for 49 different weightings of the loops. The problem with this approach is that the choice of the groups and/or the weights is rather arbitrary.

The method used to reduce the dimension of the data must preserve the information in the data. In order to develop a systematic reduction of the dimension of the data, we must define the information that we are attempting to retain and preferably provide some way to quantify this information. We will do this by using the reduced data (i.e. three or four scores for each system) to reconstruct an approximation to the original 72 loop runs. The quality of the reduction can then be measured by the discrepancy between the original and reconstructed data. Also, we can determine how well the reconstructed data retain the ranking of the systems on the segments; that is, if one system is faster than another in the original data it should also be faster in the approximation.

The reduction of the dimension of the data can be obtained from a principal component analysis of the data matrix A , that is the $m \times 72$ matrix of megaflops rates, where m is the number of systems. In principal component analysis, the data matrix A is approximated by the product BC^t where B has dimension $m \times q$ and C has dimension $72 \times q$. The B matrix contains q scores for each of the m systems. The q columns of C are the eigenvectors of $A^t A$ corresponding to the q largest eigenvalues of $A^t A$. These q eigenvalues are the squares of the singular values of A . The q elements b_{ij} of B for $1 \leq j \leq q$ then characterize the i^{th} system. The quality of this characterization is determined by how well the A matrix is approximated by the product BC^t .

The dimension of the reduced data, that is q , can be related to the predictive capability of the LFK data. If this data, for a collection of m systems, adequately describes these systems, then any given computer code could be modeled as a combination of the 72

segments weighted in some way. That is, the running time of the given code could be well estimated by computing the weighted sum of the running times of the segments. If there are m systems and n of these codes, then the running times of these codes form an $m \times n$ matrix which we denote by \mathbf{F} . Our problem is to find a $72 \times n$ matrix \mathbf{W} of weights which will predict the running time of these codes from that of the LFK loops. The matrix \mathbf{W} can be defined as the least squares solution of the equation $\mathbf{AW} = \mathbf{F}$. As we will see, the singular values of the LFK data matrix \mathbf{A} drop off very rapidly. This means that the least squares solution is not well determined. Only the first three or four weights for any given system are well determined. In this sense the dimensionality of the LFK data is quite small, certainly far less than 24. We will devote considerable attention to the selection of a reasonable value for the dimension q . However, we are unable to give a precise value for this dimension - it seems to lie between 3 and 5.

This reduction by means of the principal component analysis has the disadvantage that the scores in the matrix \mathbf{B} , for a given system are not determined solely by the benchmark times for that system. If a new system is added to the set, then the characterizations for all the systems may change. Therefore, we will discuss a second technique to define the scores for the systems.

The technique is *cluster analysis*. The 72 loop runs are decomposed into a few non-overlapping clusters. We have experimented with values of q (the number of clusters) between two and five. The cluster procedure seems to divide the loop runs in accordance with the degree of their vectorization on the vector systems. Given the clusters, then the geometric mean (or other means, see Smith[1988]) of the megaflops rates of a given system over each cluster is used to define the scores for that system. Thus, given q clusters, there are q scores for each system. Once the matrix \mathbf{B} of scores is defined, an approximation of the original data matrix is constructed from the score matrix. The quality of this approximation can then be evaluated.

The paper is organized as follows. Section 2 describes the summary statistics for both

systems and loops. Section 3 gives a mathematical description of the data reduction technique used. In sections 4 and 5 we discuss the results of our data reduction based on the principal component and cluster analyses respectively. The final section summarizes our findings and identifies directions for further work.

2 The Data and Descriptive Statistics

The data we use consist of 72 loop rates (megaflops) for 48 machine/compiler systems. We identify each loop run by a three-digit number; the first digit is the ID for the loop length and the next two-digits the loop segment number. For example, 214 designates loop 14 using the second loop length. The summary statistics for loops are listed in Table 1. It indicates that most of the loop distributions are skewed to the left (mean larger than the median). The range statistic (difference between maximum and minimum rates) can be used to identify loops that deliver high megaflops rates. The 72 loop runs are all positively correlated. The correlation coefficients for the loop runs at the first set of loop length range from 0.2309 (between loops 4 and 22) to 0.9910 (between loops 8 and 18). The correlation matrix is displayed in Table 2.

The summary statistics for different machine/compiler systems are listed in Table 3, which includes the three different means, performance ranges, and standard deviations.

3 The Principal Components Technique for Data Reduction

In this section we describe the use of principal components as a general data reduction technique. Given a data matrix \mathbf{A} of dimension $m \times n$, our goal is to find a matrix \mathbf{C} of size $n \times q$ ($q \ll n$), such that $\mathbf{B} = \mathbf{AC}$ preserves most of the information in \mathbf{A} . In our case $n = 72$. More precisely, \mathbf{C} is chosen so that the matrix $\mathbf{B} = \mathbf{AC}$ is the best linear

predictor of \mathbf{A} on the basis of q linear functions. If \mathbf{B} can replace \mathbf{A} without much loss of information, then only a small number (q) of derived variables is needed to retain most of the variation present in all of the original variables. This dimension-reducing process may aid in the interpretation of the data.

The criterion for determining the matrix \mathbf{C} is based on how well matrix \mathbf{B} can predict matrix \mathbf{A} . To reproduce \mathbf{A} from \mathbf{B} , we attempt to find a matrix \mathbf{R} of size $q \times n$ such that

$$\hat{\mathbf{A}} = \mathbf{BR}$$

is a good approximation of \mathbf{A} . Then we will have

$$\mathbf{A} = \mathbf{BR} + \epsilon = \hat{\mathbf{A}} + \epsilon.$$

The usual least squares estimator for \mathbf{R} is $(\mathbf{B}^t\mathbf{B})^{-1}\mathbf{B}^t\mathbf{A}$, or

$$\hat{\mathbf{A}} = \mathbf{AC}(\mathbf{C}^t\mathbf{A}^t\mathbf{AC})^{-1}\mathbf{C}^t\mathbf{A}^t\mathbf{A}$$

Thus, our goal is to find a matrix \mathbf{C} such that

$$\|\mathbf{A} - \hat{\mathbf{A}}\| = \|\mathbf{A} - \mathbf{AC}(\mathbf{C}^t\mathbf{A}^t\mathbf{AC})^{-1}\mathbf{C}^t\mathbf{A}^t\mathbf{A}\| \quad (4.1)$$

is a minimum.

There may not be a unique matrix \mathbf{C} which yields a minimum norm in equation 4.1. We will show that one solution for the matrix \mathbf{C} is an $n \times q$ matrix whose columns are the first q principal components of \mathbf{A} ; that is, if $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_q$ are the q largest eigenvalues of $\mathbf{A}^t\mathbf{A}$ and $\mathbf{p}_1, \dots, \mathbf{p}_q$ are the eigenvectors of unit norm corresponding to $\lambda_1, \dots, \lambda_q$ respectively, then $\mathbf{C} = (\mathbf{p}_1 \ \mathbf{p}_2 \ \dots \ \mathbf{p}_q)$.

By the spectral decomposition of \mathbf{A} , we can write

$$\begin{aligned} \mathbf{C}^t\mathbf{A}^t\mathbf{AC} &= \begin{pmatrix} \mathbf{p}_1^t \\ \vdots \\ \mathbf{p}_q^t \end{pmatrix} (\lambda_1\mathbf{p}_1\mathbf{p}_1^t + \dots + \lambda_q\mathbf{p}_q\mathbf{p}_q^t) (\mathbf{p}_1 \ \dots \ \mathbf{p}_q) \\ &= \text{diag}(\lambda_1, \dots, \lambda_q), \end{aligned}$$

since

$$\begin{aligned} \mathbf{p}_i^t \mathbf{p}_j &= 0, \text{ if } i \neq j; \\ &= 1, \text{ if } i = j. \end{aligned}$$

We have

$$\begin{aligned} \hat{\mathbf{A}} &= \mathbf{A} \mathbf{C} \text{diag}(\lambda_1^{-1}, \dots, \lambda_q^{-1}) \mathbf{C}^t \mathbf{A}^t \mathbf{A} \\ &= \mathbf{A} (\mathbf{p}_1 \cdots \mathbf{p}_q) \text{diag}(\lambda_1^{-1}, \dots, \lambda_q^{-1}) \begin{pmatrix} \mathbf{p}_1^t \\ \vdots \\ \mathbf{p}_q^t \end{pmatrix} (\lambda_1 \mathbf{p}_1 \mathbf{p}_1^t + \cdots + \lambda_q \mathbf{p}_q \mathbf{p}_q^t) \\ &= \mathbf{A} (\mathbf{p}_1 \mathbf{p}_1^t + \cdots + \mathbf{p}_q \mathbf{p}_q^t). \end{aligned}$$

A matrix result (see Rao[1973] p.70) states that the matrix $\hat{\mathbf{A}}$ of size $m \times n$ of rank q , for which $\|\mathbf{A} - \hat{\mathbf{A}}\|$ is minimum, is given by $\mathbf{A}(\mathbf{p}_1 \mathbf{p}_1^t + \cdots + \mathbf{p}_q \mathbf{p}_q^t)$; where $\mathbf{p}_1, \dots, \mathbf{p}_q$ are the first q eigenvectors of matrix $\mathbf{A}^t \mathbf{A}$, corresponding to the q largest eigenvalues of $\mathbf{A}^t \mathbf{A}$. Thus, the matrix \mathbf{C} whose columns are the first q principal components yields the minimum of $\|\mathbf{A} - \hat{\mathbf{A}}\|$.

There are other optimal properties regarding principal components besides minimizing $\|\mathbf{A} - \hat{\mathbf{A}}\|$. For more detailed discussions, see Jolliffe[1986]. The goodness of the data reduction is measured by $\|\mathbf{A} - \hat{\mathbf{A}}\|$. Since there is a large variability of the megaflops rates between and within systems, it makes more sense to look at the relative change; that is, $(a_{ij} - \hat{a}_{ij})/a_{ij}$, not $a_{ij} - \hat{a}_{ij}$. This motivates the need of a logarithmic transformation for the original data. Several other reasons also call for the log transformation. The transformation helps to correct the skewness of the loop distributions. The transformation results in data being described by geometric rather than arithmetic means which is more appropriate (Fleming and Wallace[1986]) for data which have such a wide range. In addition, it also resolves the dilemma of whether to use the megaflops rate or the time/megaflop as the unit of the measurement in the analysis (Smith[1988]).

4 Principal Component Analysis of Benchmark Data

The first 7 eigenvalues of matrix $\mathbf{A}^t\mathbf{A}$ (after log transformation) are given below.

	1	2	3	4	5	6	7
Eigenvalue	318.714	17.4104	4.72784	1.67885	1.39222	0.82538	0.60049
Proportion	0.9163	0.0501	0.0136	0.0048	0.0040	0.0024	0.0017
Cumulative	0.9163	0.9663	0.9799	0.9848	0.9888	0.9911	0.9929

Here the k^{th} "Proportion" entry is the ratio of the k^{th} eigenvalue to the trace of $\mathbf{A}^t\mathbf{A}$, and the "Cumulative" entry the ratio of the sum of the first k eigenvalues to the trace. Since the sum of eigenvalues is the trace of $\mathbf{A}^t\mathbf{A}$, i.e. the total sum of squares of \mathbf{A} , the ratio of each eigenvalue to the sum can be viewed as the proportion of the total sum of squares accounted for by the corresponding component.

The values of $\|\mathbf{A} - \hat{\mathbf{A}}\|$ resulting from using one to seven components are given below, which indicates that the values of $\|\mathbf{A} - \hat{\mathbf{A}}\|$ begin to level off after 3 or 4 components have been extracted.

No. of Components	1	2	3	4	5	6	7
$\ \mathbf{A} - \hat{\mathbf{A}}\ $	37.38	23.70	18.30	15.95	13.69	12.16	10.91

If we consider only the reduction in the size of the eigenvalues of $\mathbf{A}^t\mathbf{A}$ and the difference $\|\mathbf{A} - \hat{\mathbf{A}}\|$, then it is difficult to decide how many components should be used to approximate the data matrix \mathbf{A} . However, an inspection of the elements of the \mathbf{C} matrix shows a correlation with the nature of the code segments. The j^{th} component, shown below, is simply the j^{th} column of the \mathbf{C} matrix. The first component clearly measures overall performance of systems as would be expected since all the correlations between the 72 loop runs are positive and this component accounts for 91.63% of the total variation.

Run	Component 1	Run	Component 1	Run	Component 1
101	0.14418	201	0.15974	301	0.16854
102	0.08856	202	0.11590	302	0.11592
103	0.11448	203	0.13637	303	0.15791
104	0.07603	204	0.10066	304	0.13150
105	0.09109	205	0.09168	305	0.09192
106	0.07458	206	0.09478	306	0.10090
107	0.15387	207	0.17048	307	0.17601
108	0.14486	208	0.16454	308	0.16442
109	0.14713	209	0.16760	309	0.16765
110	0.11838	210	0.13147	310	0.13194
111	0.07883	211	0.08144	311	0.08280
112	0.11333	212	0.13226	312	0.14183
113	0.07514	213	0.08383	313	0.08504
114	0.09635	214	0.09908	314	0.09947
115	0.07681	215	0.07890	315	0.07881
116	0.07300	216	0.07194	316	0.07209
117	0.10019	217	0.09925	317	0.09941
118	0.13859	218	0.16016	318	0.16010
119	0.09861	219	0.10117	319	0.10123
120	0.10427	220	0.10418	320	0.10397
121	0.12303	221	0.13286	321	0.13659
122	0.11467	222	0.12890	322	0.12892
123	0.10885	223	0.11174	323	0.11171
124	0.06075	224	0.06937	324	0.07914

We observe that the faster loops tend to have a negative second component. That is, the second component, given in the table below, contrasts loops that deliver high megaflops

rates (negative coefficients) with the rest (positive coefficients). Also, we observe that the degree of the loop vectorization is correlated with the magnitude of the second component. For example, the sign of the coefficients of runs 107, 207 and 307 indicates that the loop 7 is a vectorized loop; and the magnitude of the coefficients reveals that the loop performance increases as the loop length increases. Similarly, the loop 9 is also a vectorized loop, however, the megaflops rate tops out at the second segment of loop length and further increase of loop length will not increase the performance as indicated by the magnitude of the coefficients of runs 209 and 309. Thus, after overall performance has been accounted for, the next source of variation is between systems with vectorization capability and systems without that capability.

Run	Component 2	Run	Component 2	Run	Component 2
101	-0.09900	201	-0.15848	301	-0.18935
102	0.14579	202	0.03460	302	0.03422
103	0.02971	203	-0.06082	303	-0.14264
104	0.21104	204	0.11288	304	-0.00798
105	0.12086	205	0.11992	305	0.11773
106	0.18601	206	0.10756	306	0.08142
107	-0.13936	207	-0.19955	307	-0.21774
108	-0.06432	208	-0.13863	308	-0.13815
109	-0.09708	209	-0.17267	309	-0.17254
110	0.01296	210	-0.04306	310	-0.04044
111	0.17253	211	0.16245	311	0.15348
112	0.07387	212	0.00585	312	-0.02722
113	0.23655	213	0.20599	313	0.20063
114	0.12899	214	0.11779	314	0.11658
115	0.11636	215	0.09972	315	0.09886
116	0.15707	216	0.16114	316	0.16001
117	0.06824	217	0.07173	317	0.07185
118	-0.04980	218	-0.12188	318	-0.12138
119	0.08271	219	0.06056	319	0.06047
120	0.02913	220	0.02910	320	0.02987
121	0.01883	221	0.00072	321	-0.00521
122	-0.00758	222	-0.05586	322	-0.05523
123	0.02497	223	0.01623	323	0.01629
124	0.16801	224	0.13680	324	0.09621

The third component, given below, identifies the loop length effect for each vectorized loop. The smaller the coefficient, the larger the impact of the loop lengths to the loop

performance. For example, the performance of the loop 1 is an increasing function of all the 3 loop lengths, while for loop 2, it's an increasing function only for the first 2 loop lengths. The large and almost identical coefficients in loop 20 indicate that the loop is scalar and therefore there is no loop length effect. Overall, the first 3 components account for a substantial proportion of the total variation (97.99%).

Run	Component 3	Run	Component 3	Run	Component 3
101	0.03088	201	-0.03163	301	-0.09567
102	0.13383	202	0.05181	302	0.05461
103	0.04943	203	-0.04193	303	-0.16055
104	0.10000	204	0.03123	304	-0.10619
105	0.08464	205	0.08652	305	0.08539
106	0.07104	206	0.01071	306	-0.01658
107	0.11108	207	0.02954	307	-0.00496
108	0.11022	208	0.00144	308	0.00079
109	0.10966	209	0.01088	309	0.01084
110	0.03193	210	-0.04147	310	-0.04184
111	0.01063	211	0.01121	311	0.00500
112	-0.011278	212	-0.19050	312	-0.25377
113	-0.14047	213	-0.19008	313	-0.19405
114	-0.06348	214	-0.07168	314	-0.07946
115	0.01566	215	0.00653	315	0.00664
116	0.01896	216	0.01970	316	0.01901
117	0.15355	217	0.15349	317	0.15349
118	0.06302	218	-0.03120	318	-0.03116
119	0.15346	219	0.17395	319	0.17432
120	0.17536	220	0.17653	320	0.17329
121	-0.04600	221	-0.10913	321	-0.13565
122	-0.08217	222	-0.13985	322	-0.13970
123	0.18905	223	0.17340	323	0.17230
124	-0.15600	224	-0.23161	324	-0.33910

4.1 Goodness of Reduction

Although the principal components were derived solely based on the minimization of $\|A - \hat{A}\|$, other criteria can also be used to assess the adequacy of the dimension reduction. Our objective is the simplification of the benchmark data. This requires a reduction in its dimension. However, looking only at the norm $\|A - \hat{A}\|$ and the eigenvalues of $A^t A$ does not provide a clear indication of the number of components which should be used to reconstruct \hat{A} . We will use three extra measures of the quality of the approximation of A by \hat{A} . The first is a comparison of the geometric means of the megaflops rates for each system obtained from the original data A with the means obtained from the approximate data of \hat{A} . The second is the performance range, the difference between the maximum and minimum megaflops rates, for each system. The third is the Spearman rank correlation coefficient (Noether[1967]) for the 72 loop runs for each system. The Spearman correlation indicates how well the rank ordering of the loop rates is preserved. Specifically, for each system, the Spearman correlation coefficient between the ranks of the 72 megaflops rates of A and \hat{A} is calculated. A large Spearman correlation indicates a good preservation of the ranking. Table 4 lists the results for one to four components (used in the data reduction). This table shows the ratio of the geometric mean of the reconstructed data to the geometric mean of the original data, and also gives a similar ratio for the range of the data for each system.

Table 4 shows that the Spearman coefficient tends to be smaller for the scalar systems because there is much less variation of the loop rates compared to that for a vector system. Therefore the rank order is not as well defined for the scalar systems so that the data in the \hat{A} must be more accurate in order to preserve the rank order.

Note that the geometric mean is already quite accurate when only one component is used to construct \hat{A} . Also, in Table 4, there is little change in the geometric means obtained from \hat{A} in going from one to four components.

However, to preserve the rank ordering of the loop rates for each system, as measured

by the Spearman correlation, requires three components. There is considerable change in the Spearman correlations in going from one to two components, and some change from two to three components. Table 4 is concerned with the ranking of the loops for each system; thus there are 72 elements in the ranking with the fastest loop at the head of the ranking. Perhaps a ranking of the systems for each loop is of greater significance to benchmarking. Here there are 48 systems in the ranking for each loop with the fastest system at the head of the ranking. For such a ranking of the systems on each loop, the Spearman correlations are sufficiently large (i.e. generally greater than 0.9) when only two components are used to generate \hat{A} .

It is very difficult to accurately reproduce the range using only a few components to construct the \hat{A} . The actual range for the ETA205-V is 167 and the approximate range using four components is only 56. For other vector systems the error was not this large, but was still in the 20% region. Approximating extreme values is very difficult. It is difficult to determine how important these extreme values are in the evaluation of the system. Nevertheless, it seems that three components are the minimum number required to adequately preserve the central tendency, the spread, and the rank ordering of the original data.

Another common practice of assessing the goodness-of-fit is to examine the residuals. Plots of elements of $A - \hat{A}$ are displayed in Figures 1a, 1b and 1c. They show that, among 72 runs, loop 24's have the most extreme outlying points. Detailed examination reveals that most of the outliers are from the same group of machines. For example, the top 3 outliers of runs 224 and 324 are Amdahl (1400VP-V, 1200VP-V, and 1500VP-V); the top outlier of loop 23 is Apollo300-32; the top 2 outliers of runs 121, 219, and 319 are Convex (V-32 and V-64); the top outlier of loop 15 is Alliant-V-64-P. In general, most of the residuals lie within the band of $(-0.5, 0.5)$ indicating that 3 principal components are probably sufficient in summarizing the data.

4.2 Component Scores

The matrix $\mathbf{B} = \mathbf{AC}$ is the matrix of summary scores obtained from the reduction. We also point out here that the matrix \mathbf{C} (and consequently \mathbf{B}) is not uniquely determined and that any other matrix $\tilde{\mathbf{C}}$ of size $n \times q$ whose columns span the same space as the columns of \mathbf{C} will also solve the data reduction problem, i.e. $\tilde{\mathbf{C}} = \mathbf{CP}$ where \mathbf{P} is any $q \times q$ nonsingular matrix is also an optimal solution. It is easily shown that there exists a 3×3 nonsingular matrix \mathbf{P} such that the first column of the matrix $\tilde{\mathbf{C}} = \mathbf{CP}$ will sum to 1, and the negative and the positive coefficients of the second and the third columns of $\tilde{\mathbf{C}}$ will sum to -1 and 1 respectively. This transformation will allow us to interpret the scores as geometric means. Let \mathbf{A} and \mathbf{B} be the data and score matrices in log scale and $\tilde{\mathbf{A}}$ the original data matrix, i.e. $\mathbf{A} = \log(\tilde{\mathbf{A}})$, then $\tilde{\mathbf{B}} = \mathbf{A}\tilde{\mathbf{C}}$. The rows of $\tilde{\mathbf{B}}$ now represent the summary scores corresponding to each of the 48 systems.

For score 1, we have

$$\begin{aligned}\tilde{b}_{i1} &= \sum_{k=1}^{72} a_{ik} \tilde{c}_{k1} \\ &= \sum_{k=1}^{72} \tilde{c}_{k1} \log \tilde{a}_{ik} \\ &= \log\left(\prod_{k=1}^{72} \tilde{a}_{ik}^{\tilde{c}_{k1}}\right).\end{aligned}$$

If we let $\tilde{b}_{ij} = \log \tilde{b}_{ij}$, which is the j^{th} summary score for system i in the original scale, we have

$$\tilde{b}_{i1} = \prod_{k=1}^{72} \tilde{a}_{ik}^{\tilde{c}_{k1}},$$

which is a weighted geometric mean of the 72 loop rates with weights \tilde{c}_{k1} , $k = 1, 2, \dots, 72$.

For scores 2 and 3 (i.e. $j = 2$ and 3), we have

$$\begin{aligned}\tilde{b}_{ij} &= \sum_{\{\tilde{c}_{kj} > 0\}} a_{ik} \tilde{c}_{kj} + \sum_{\{\tilde{c}_{kj} < 0\}} a_{ik} \tilde{c}_{kj} \\ &= \log\left(\prod_{\{\tilde{c}_{kj} > 0\}} \tilde{a}_{ik}^{\tilde{c}_{kj}} / \prod_{\{\tilde{c}_{kj} < 0\}} \tilde{a}_{ik}^{|\tilde{c}_{kj}|}\right),\end{aligned}$$

or

$$\bar{b}_{ij} = \prod_{\{\bar{c}_{kj} > 0\}} \bar{a}_{ik}^{\bar{c}_{kj}} / \prod_{\{\bar{c}_{kj} < 0\}} \bar{a}_{ik}^{|\bar{c}_{kj}|},$$

which is a ratio of 2 weighted geometric means. These scores are tabulated in Table 5.

Score 1 measures the overall performance of each system. For vector systems, score 1 can be significantly larger than the geometric mean, since the weights \bar{c}_{k1} have slightly larger value on vectorizable loops. For scalar systems, score 1 and the geometric mean are very close, implying that vectorizable loops play no significant roles here. Score 2 is the ratio of the scalar performance to the vector performance and can be used to easily identify the vector systems and their vectorizability. The smaller the value of score 2, the more vectorizable the system. Score 3 measures the loop length effect for the vector systems. Again, the small value of score 3 implies the significant length effect.

Another advantage of reducing the dimensionality of the data is that we are able to plot the data. The original 72-dimensional data are impossible to visualize. The first 3 components give the *best-fitting* 3-dimensional subspace and preserve a substantial proportion of the total variation. Figure 2 gives the plot of the machines with respect to the first 2 scores. The size of the markers in the plot is proportional to the reciprocal of the third score. The vector and scalar systems are in separate clusters, with the vector systems at the bottom of the plot. Note that when the Cray systems are run in scalar mode, they appear in the scalar cluster. The markers are: Alliant -“⊏”, Amdahl -“△”, Apollo -“⊕”, Convex -“▽”, Cray -“+”, DEC -“*”, ETA -“●”, IBM -“o”, NEC -“■”, SCS - “◇”, Sperry - “×”, and all others - “⊗”.

5 Cluster Analysis on Loop Runs

As we mentioned earlier, although the principal components have some desired optimal properties in data reduction, there is one disadvantage of being data dependent. If a new system is added to the benchmark data set, then the scores for all the systems may change. In order to eliminate this data dependence, the 72 loop runs are divided into q clusters.

The geometric mean of the megaflops rates over each cluster is used to define q scores for each system. Once the clusters are defined, then the scores for each system are completely independent of the scores for the other systems. However, the clusters must be defined so that these scores give a good characterization of the systems.

The decomposition of the loop runs is obtained from a $72 \times q$ matrix G of weights. This matrix is used to generate a score matrix B in the same way that the matrix C of principal components generates a score matrix, that is $B = AG$. This matrix G is restricted to have a single non-zero element in each row. This non-zero element identifies the cluster membership and the weight within the cluster for the loop run. Therefore G can be used to decompose the loop runs into q clusters. The elements of G must be chosen so that the score matrix B is the best possible predictor of the original data matrix A . In fact, we could formulate the problem by the same technique used in section 3, i.e. our goal would be to find a matrix $G \in \Gamma$, where Γ is the collection of all the $72 \times q$ matrices having only a single non-zero element in each row, such that

$$\|A - AG(G^t A^t AG)^{-1} G^t A^t A\| \quad (6.1)$$

is a minimum.

The minimization of (6.1) is a difficult (computational) problem. However, in the general case (i.e. no restriction imposed on G) the problem is equivalent to the minimization of the trace of the residual variance of predicting A based on the linear predictor AG (see Rao[1973], p.593). This residual variance can be expressed as the covariance matrix of A given AG (that is, the covariance conditional on AG), which is

$$\Sigma - \Sigma G(G^t \Sigma G)^{-1} G^t \Sigma, \quad (6.2)$$

where Σ , of size 72×72 , is the covariance matrix of the loop runs. To minimize the trace of (6.2), we need to maximize $\text{trace}(\Sigma G(G^t \Sigma G)^{-1} G^t \Sigma)$. If there were no constraint imposed on G , the optimum choice of G would be $G = C$, the q largest principal components of

Σ . Also, since (see Rao[1973] p.592) for any $72 \times q$ matrix X

$$\begin{aligned} \max_X \text{trace}(\Sigma X (X^t \Sigma X)^{-1} X^t \Sigma) &= \text{trace}(\Sigma C (C^t \Sigma C)^{-1} C^t \Sigma) \\ &= \text{trace}(C^t \Sigma C), \end{aligned}$$

this motivates us to reformulate the problem; instead of finding $G \in \Gamma$ to minimize (6.1), we will find $G \in \Gamma$ such that

$$\text{trace}(G^t \Sigma G) = g_1^t \Sigma g_1 + \cdots + g_q^t \Sigma g_q \quad (6.3)$$

is a maximum, where g_i is the i^{th} column of G . If we denote by h_i the column vector containing the non-zero elements of g_i , and Σ_i the covariance submatrix of Σ corresponding to these non-zero elements, then $g_i^t \Sigma g_i = h_i^t \Sigma_i h_i$. In addition,

$$\max_{h_i} h_i^t \Sigma_i h_i = c_i^t \Sigma_i c_i,$$

where c_i is the eigenvector corresponding to the largest eigenvalue of Σ_i . Thus, the elements of G which maximize (6.3) can be easily determined if the cluster structure is known.

In this paper, we employ the VARCLUS procedure of the SAS[1986] to find the cluster components G . It begins with all loops in one cluster and repeats the following steps until q clusters are obtained.

1. The principal components for each cluster are computed, that is, the eigenvectors of each Σ_i . The cluster having the largest second eigenvalue is chosen for further splitting.
2. The chosen cluster is split into two clusters by finding the first two principal components, performing a rotation (Harman[1976]) and assigning each loop run to the rotated (cluster) component with which it has the higher squared correlation.

Once q clusters are obtained then an iterative procedure is used to reassign loop runs to clusters in order to maximize the trace in (6.3).

Since the principal components were obtained without the constraint, a given number of cluster components does not explain as much variance as the same number of principal components. However, the cluster components are easier to interpret than the principal components.

The cluster results obtained from the VARCLUS procedure are given in the tables below. For the two-cluster case, the elements of the matrix G , multiplied by 10^6 , are given in the following table.

Run	Cluster 1	Cluster 2	Run	Cluster 1	Cluster 2	Run	Cluster 1	Cluster 2
101		33061	201		33164	301		32941
102	25069		202	25073		302	25094	
103	25102		203		33011	303		32840
104	24585		204	25272		304		32541
105	25440		205	25425		305	25438	
106	24978		206	25059		306	24753	
107		33129	207		33231	307		33075
108		32830	208		33204	308		33205
109		32947	209		33220	309		33222
110		32580	210		33001	310		33006
111	25129		211	25014		311	24981	
112		32792	212		32997	312		32829
113	24702		213	24742		313	24741	
114	25129		214	25083		314	25056	
115	24804		215	24740		315	24721	
116	25292		216	25251		316	25262	
117	25394		217	25377		317	25379	
118		32941	218		33300	318		33298
119	25159		219	24654		319	24652	
120	25262		220	25255		320	25252	
121		32726	221		33043	321		33009
122		32737	222		32753	322		32748
123	25063		223	25090		323	25085	
124	23645		224	22301		324		27334

As we can see, the weights within each cluster are nearly constant for the above case. In fact, as the number of clusters increases, the weights tend to be even less variable within

each cluster because the clusters become more homogeneous. Thus, we can treat the loop runs within each cluster equally without loss of much information. The clusters obtained from the VARCLUS procedure for $q = 3, 4$, and 5 are shown below.

Cluster 1	Cluster 2	Cluster 3
102 202 302	101 201 301	124 224 324
103	203 303	
104 204	304	
105 205 305	107 207 307	
106 206 306	108 208 308	
111 211 311	109 209 309	
113 213 313	110 210 310	
114 214 314	112 212 312	
115 215 315	118 218 318	
116 216 316	121 221 321	
117 217 317	122 222 322	
119 219 319		
120 220 320		
123 223 323		

Cluster 1	Cluster 2	Cluster 3	Cluster 4
102	101 201 301	124 224 324	202 302
104 204	203 303		103
105 205 305	304		206 306
106	107 207 307		113 213 313
111 211 311	108 208 308		114 214 314
116 216 316	109 209 309		115 215 315
117 217 317	110 210 310		
119 219 319	112 212 312		
120 220 320	118 218 318		
123 223 323	121 221 321		
	122 222 322		

Cluster 1	Cluster 2	Cluster 3	Cluster 4	Cluster 5
102	201 301	124 224 324	202 302	101
104 204	203 303		103	304
105 205 305	207 307		206 306	107
106	209 309		113 213 313	108 208 308
111 211 311	212 312		114 214 314	109 209 309
116 216 316	218 318		115 215 315	110 210 310
117 217 317	222 322			112
119 219 319				118
120 220 320				121 221 321
123 223 323				122

The cluster analysis apparently groups the loops according to ease of vectorization and megaflops rate. Consider the situation when the loops are broken into four clusters. The first cluster consists generally of loops that do not vectorize, the second of loops which

vectorize very well and give high megaflops rates, the third cluster is an anomaly which consists of a single loop, and the fourth consists of loops which vectorize only moderately well and do not give high megaflops rates. This can be seen from Table 6. This table lists, for each of the 72 loop runs, the first and third quartiles (Q1 and Q3) of the megaflops rates over the 48 systems. Also listed is a vectorization figure. The report by McMahon gives the extent of vectorization for each of the 24 loops for 6 systems: Cray-1, Fujitsu, Cyber205, Convex C1, NECSX-2, and IBM3090. Table 5 of McMahon's report shows full, partial, or no vectorization for each of these systems. The "vector" column in Table 6 gives this number for each loop. If the vectorization number is 6, then all systems fully vectorized the loop. If this number is 4p, then four systems partially vectorized the loop and two systems gave no vectorization. It is clear that most of the loops in cluster 1 did not vectorize on any systems, whereas most loops in cluster 2 vectorized on all systems. The average loop in cluster 4 vectorized on only three systems, thus this cluster is intermediate between 1 and 2. Cluster 3 contains only loop 24 which is an anomalous case. This loop is the following:

```

M = 1
DO 24 K = 2, M
    IF (X(K) .LT. X(M)) M=K
24 CONTINUE

```

The Amdahl vector systems ran this loop an order of magnitude, or more, faster than the other systems. Therefore the results for this loop have a different structure than those for the other loops; so much different that this loop forms a cluster by itself. Perhaps the Amdahl system has a hardware instruction to locate the smallest element in an array, and the compiler is clever enough to generate that instruction. At any rate, this loop seems to be an anomaly. It is rather remarkable that this cluster analysis seems to select the loops based on vectorization.

Next we consider a method to define scores for each system based on this decompo-

sition of the loops into clusters. Given a decomposition of the loops into q clusters, the corresponding scores are defined as the geometric mean of the megaflops rates for the given system over the clusters. Thus, if there are m systems, then we have defined an $m \times q$ score matrix \mathbf{B} . This matrix is shown in Tables 7 for $q = 2$ and $q = 4$. For the case of four clusters, the first score is the geometric mean over loops which vectorize poorly, the second over loops which vectorize very well, the third over the single loop which finds the smallest element in an array, and the fourth over loops which are partially vectorized.

From the score matrix \mathbf{B} , we construct an approximation $\hat{\mathbf{A}}$ of the original data matrix \mathbf{A} . The approximation is obtained by least squares. The values of $\|\mathbf{A} - \hat{\mathbf{A}}\|$ resulting from using two to five clusters are given below.

No. of Clusters	2	3	4	5
$\ \mathbf{A} - \hat{\mathbf{A}}\ $	24.61	20.58	17.32	16.30

For a given value of q , the L_2 norm $\|\mathbf{A} - \hat{\mathbf{A}}\|$ based on the principal components is smaller, as we might expect, since the principal component approximation is optimal. Also, from this approximation $\hat{\mathbf{A}}$ the geometric mean and range of the loop runs for each system can be computed. These are shown in Table 8. The mean and range can be compared with those obtained from the original matrix which are displayed in Table 3. In addition, the mean and range can be compared with those in Table 4 obtained from the approximation based on the principal components. The approximation based on cluster analysis requires four clusters to give roughly the same accuracy for the geometric mean and Spearman rank correlation as three components of the principal component analysis. However, the estimate of the range obtained from the scores based on the cluster analysis is superior to that obtained from the scores based on the principal components.

6 Concluding Remarks

In this paper, we have investigated the “dimensionality” of the 24 Livermore loops. In this context, dimension is defined as the number of linear combinations of the loop timings

that can be used as “scores” to characterize a computer hardware/software system. This dimension is based on a singular value decomposition of the loop timings over a set of 48 computer systems. Therefore, the dimension is not well defined, since it is difficult to determine when a small singular value should be set to zero and the rank of the data matrix reduced. However, the dimension is certainly greater than one; a single number, such as the Linpack timing, has too little predictive value. We find that three to five of these scores are required to reconstruct the original Livermore benchmark data fairly accurately.

We also present two methods to define the scores for the systems. The first is optimal in a certain sense and is based on a principal component analysis. It has the disadvantage that the interpretation of the scores is not obvious. The second method uses a grouping of the loops into clusters. The scores for a given system are the geometric means of the megaflops rates taken over each cluster. These clusters are closely related to the vectorization of the loops.

There are other ways to approach the data reduction problem. For instance one could ask if a subset of the 72 loop runs will provide essentially the same information as the full set. This question could be addressed by performing a *best subset analysis* on the LFK data. Another important issue is the external validation of the scores derived in this paper. In particular, the predictive power of the loops could be tested by running the loops on a set of 10 to 15 systems along with a few small “production” codes. Then the scores obtained from the loops could be used to give a least squares prediction of the running times on the production codes. If the loops and the resulting scores really characterize the systems, the prediction should be fairly good.

References

- [1] Fleming, P. J. and Wallace, J. J. (1986), "How Not to Lie with Statistics: The Correct Way to Summarize Benchmark Results", *Communications of the ACM*, vol. 29, no. 3.
- [2] Harman, H. H. (1976), *Modern Factor Analysis, 3rd Edition*, Chicago: University of Chicago Press.
- [3] Jolliffe, I. T. (1986), *Principal Component Analysis*, New York: Springer-Verlag.
- [4] Kuck, D. J. and Sameh, A. H. (1987), "A Supercomputing Performance Evaluation Plan", Supercomputer Conference. Athens, Greece.
- [5] Noether, G. E. (1967), *Elements of Nonparametric Statistics*, New York: John Wiley.
- [6] McMahon, F. H. (1986), "The Livermore Fortran Kernels: A Computer Test of the Numerical Performance Range", *Reprot UCRL - 537415*, Lawrence Livermore National Laboratory.
- [7] Rao, C. R. (1973), *Linear Statistical Inference and Its Applications*, New York: John Wiley and Sons.
- [8] SAS Institute Inc. (1986), *SAS User's Guide: Statistics*, Gary, N.C.
- [9] Smith, J. E. (1988), "Characterizing Computer Performance with A Single Number", *Communications of the ACM*, vol. 31, no. 10.

Table 1. Statistics of Loop Runs

Loop	Mean	Stddev	Min	Med	Max	Loop	Mean	Stddev	Min	Med	Max
101	17.64	28.80	0.007	7.60	158.51	113	1.85	1.75	0.005	1.44	7.00
201	36.57	84.93	0.007	9.02	529.75	213	2.44	3.06	0.005	1.63	14.34
301	61.87	159.17	0.007	9.02	800.05	313	2.56	3.43	0.005	1.70	16.78
102	3.32	3.32	0.007	2.26	15.60	114	3.53	3.98	0.005	2.42	19.59
202	6.80	9.15	0.007	3.83	49.94	214	3.88	4.75	0.005	2.48	24.16
302	6.81	9.16	0.007	3.77	49.94	314	3.96	5.02	0.005	2.53	25.79
103	6.66	8.09	0.007	4.15	43.91	115	2.50	2.17	0.008	2.45	8.91
203	14.30	23.91	0.007	5.42	122.01	215	2.60	2.19	0.008	2.39	8.74
303	44.50	112.01	0.007	5.75	528.67	315	2.60	2.19	0.008	2.41	8.74
104	2.53	2.94	0.006	1.54	15.65	116	2.24	2.22	0.010	1.63	9.85
204	4.41	5.48	0.006	2.60	28.70	216	2.19	2.19	0.010	1.57	9.85
304	15.48	32.87	0.006	3.24	164.18	316	2.20	2.21	0.010	1.59	9.85
105	3.47	3.31	0.007	2.50	13.17	117	4.88	4.65	0.011	3.46	18.10
205	3.54	3.38	0.007	2.52	13.68	217	4.78	4.58	0.011	3.24	17.89
305	3.55	3.38	0.007	2.56	13.58	317	4.79	4.57	0.011	3.35	17.89
106	2.27	2.19	0.005	1.67	10.74	118	12.80	16.61	0.006	7.38	66.72
206	3.65	3.74	0.005	2.47	18.74	218	33.56	71.01	0.006	8.59	349.42
306	4.46	5.35	0.005	2.86	29.30	318	33.52	70.99	0.006	8.59	349.42
107	23.65	36.39	0.009	10.65	178.95	119	4.29	3.83	0.008	3.64	16.17
207	52.53	123.12	0.009	12.02	720.82	219	4.74	4.22	0.008	4.78	18.12
307	75.11	196.46	0.009	12.38	1042.33	319	4.75	4.22	0.008	4.78	18.11
108	15.91	21.12	0.006	6.79	87.20	120	5.75	5.33	0.010	3.84	19.36
208	38.68	82.36	0.006	10.41	415.70	220	5.75	5.29	0.010	3.84	19.29
308	38.59	82.31	0.006	10.40	415.68	320	5.71	5.29	0.010	3.86	19.35
109	18.28	26.42	0.008	8.47	121.10	121	8.34	12.24	0.006	3.83	65.72
209	45.91	112.05	0.008	11.03	705.20	221	13.44	27.01	0.006	3.57	156.56
309	45.97	112.12	0.008	11.06	705.28	321	17.08	40.58	0.006	3.27	253.03
110	7.61	9.39	0.010	3.99	33.96	122	8.03	12.43	0.006	2.80	43.37
210	13.46	23.74	0.010	4.49	120.75	222	15.81	33.22	0.006	3.25	183.36
310	13.48	23.78	0.010	4.42	120.75	322	15.80	33.20	0.006	3.25	183.34
111	2.48	2.30	0.008	1.70	8.32	123	6.51	6.22	0.007	4.60	23.30
211	2.69	2.52	0.008	1.70	8.32	223	7.02	6.80	0.007	4.79	24.48
311	2.79	2.65	0.008	1.70	8.70	323	7.02	6.80	0.007	4.79	24.44
112	5.77	8.21	0.004	2.48	39.32	124	2.07	2.76	0.033	1.03	12.53
212	12.72	25.73	0.004	2.89	147.41	224	3.94	9.48	0.033	1.07	45.80
312	21.64	50.11	0.004	3.05	242.80	324	12.74	46.89	0.033	1.26	266.58

Table 2. Correlations of Loops (times 10^{-4})

	102	103	104	105	106	107	108	109	110	111	112
101	4912	8621	3098	5726	3368	9806	9297	9682	8845	5369	9590
102		6655	9384	8915	9309	5020	6185	5451	5891	8632	5234
103			5964	8041	5722	8577	8443	8216	7951	8016	8518
104				8698	9523	3205	4409	3465	4067	8585	3586
105					8767	5905	6773	6153	6847	9533	5894
106						3366	4528	3793	4312	8854	4026
107							9686	9866	9375	5354	9407
108								9764	9754	6135	9207
109									9587	5486	9299
110										6055	8776
111											5964

Table 2. Correlations of Loops (continued)

	113	114	115	116	117	118	119	120	121	122	123	124
101	7357	8785	6119	6373	6321	9234	5585	6363	9520	8996	6533	4633
102	8392	6960	9150	9272	9522	5928	9006	8817	4988	4322	9294	4284
103	8298	9396	7455	8032	7871	8438	7733	7628	9170	7193	8127	7159
104	7069	5579	8422	8849	8751	4216	8628	7848	3586	2309	8453	4546
105	8047	7752	8883	9245	9564	6788	9764	9410	6629	5204	9670	6658
106	7124	5542	8522	8846	8826	4321	8685	7884	3764	2749	8430	4375
107	7717	9086	6097	6347	6385	9694	5709	6755	9524	9492	6708	5475
108	8450	9258	6870	7014	7296	9910	6682	7731	9022	9576	7633	5818
109	7851	8916	6368	6521	6718	9762	6010	7126	9298	9576	7029	5234
110	8050	8869	6626	6679	7143	9833	6692	7861	8770	9550	7573	6088
111	7813	7435	8517	9153	9152	6043	9212	8450	6136	4458	9028	6456
112	7497	8491	6238	6447	6488	9102	5837	6454	8960	8795	6726	4944
113		9224	8224	8323	8699	8269	8018	8574	7333	7105	8831	6252
114			7500	7999	8046	9199	7502	8076	9100	8282	8220	6895
115				9115	9354	6855	8754	8893	6165	5387	9176	4655
116					9671	6878	9052	8880	6829	5376	9203	5653
117						7149	9585	9542	6651	5718	9790	5522
118							6603	7721	9118	9565	7534	6244
119								9298	6381	5079	9648	5913
120									6918	6587	9745	6399
121										8602	6942	6417
122											6076	4991
123												6104

Table 3. Statistics of Machine/Compiler Systems

System	Harmonic	Geometric	Average	Std.	Min.	Max.	Range
ALLIANT-S-32	0.637	0.721	0.813	0.390	0.303	1.580	1.277
ALLIANT-V-32	0.801	1.164	1.648	1.434	0.096	5.390	5.294
ALLIANT-S-64	0.573	0.627	0.685	0.283	0.287	1.250	0.963
ALLIANT-V-64-P	1.199	2.257	5.026	6.839	0.280	29.200	28.920
AMDAHL5890-S	6.208	7.020	7.664	2.841	1.730	11.970	10.240
AMDAHL1500VP-V	10.087	17.334	31.248	33.741	2.230	116.300	114.070
AMDAHL1200VP-V	11.730	24.712	65.528	97.816	2.700	435.520	432.820
AMDAHL1400VP-S	6.294	7.396	8.354	3.786	1.700	16.000	14.300
AMDAHL1400VP-V	11.940	27.174	88.140	154.868	2.670	819.450	816.780
APOLLO300-32	0.013	0.015	0.020	0.027	0.005	0.143	0.138
APOLLO660-32	0.101	0.109	0.115	0.040	0.044	0.225	0.181
APOLLO300-64	0.007	0.007	0.008	0.005	0.004	0.033	0.029
APOLLO660-64	0.070	0.073	0.076	0.020	0.036	0.112	0.076
SUN3-64	0.287	0.321	0.361	0.189	0.102	0.910	0.808
RIDGE32	0.196	0.202	0.208	0.049	0.121	0.292	0.171
CDC875	3.265	3.653	4.036	1.737	1.240	8.380	7.140
CYBER176	2.779	3.215	3.664	1.798	1.110	8.270	7.160
CELERITY-32	0.231	0.259	0.292	0.154	0.091	0.809	0.718
CONVEX-S-32	1.123	1.278	1.430	0.679	0.400	3.600	3.200
CONVEX-V-32	1.233	2.640	5.246	6.070	0.123	23.600	23.477
CONVEX-S-64	0.925	1.060	1.193	0.561	0.338	2.750	2.412
CONVEX-V-64	1.035	1.888	3.235	3.301	0.111	12.790	12.679
CRAY1-S	4.801	5.513	6.451	3.831	2.290	15.430	13.140
CRAY1-V	6.589	11.977	23.547	26.339	1.430	95.420	93.990
CRAYXMP-S	5.726	6.647	7.859	4.797	2.620	19.160	16.540
CRAYXMP-V	8.289	17.021	39.363	47.549	2.140	162.190	160.050
CRAYXMP-CFT-S	5.694	6.975	8.524	5.562	1.580	22.700	21.120
CRAYXMP-CFT-V	7.957	17.052	37.546	44.354	1.520	167.720	166.200
CRAY2-S	3.682	4.393	5.288	3.262	1.640	12.120	10.480
CRAY2-V	5.135	11.278	29.041	37.941	1.260	146.400	145.140
MICROVAX2	0.163	0.173	0.181	0.050	0.061	0.280	0.219
VAX8800	0.885	0.948	1.001	0.304	0.307	1.644	1.337
VAX8800-32	1.243	1.343	1.432	0.486	0.460	2.410	1.950
ELXSI6420	1.078	1.179	1.291	0.561	0.517	2.740	2.223
ETA205-S	3.366	4.323	5.570	4.266	0.880	17.600	16.720
ETA205-V	4.253	7.352	17.238	30.321	0.850	167.920	167.070
IBM3033	1.369	1.523	1.643	0.564	0.420	2.400	1.980
IBM3081	2.332	2.441	2.539	0.669	1.190	3.570	2.380
IBM3090-S	6.013	6.571	7.070	2.458	2.900	11.010	8.110
IBM3090-V	7.034	9.104	12.443	11.187	2.020	47.500	45.480
NECSX2-S	11.228	12.895	14.895	8.384	4.590	38.160	33.570
NECSX2-V	18.545	42.022	135.237	221.717	4.470	1042.330	1037.860
FPS264-64	4.690	6.009	7.573	5.315	1.230	21.640	20.410
HONEYWELDPS-90	3.615	4.394	5.340	3.297	1.530	13.570	12.040
NASXL-60	8.539	11.329	13.860	7.631	1.920	28.000	26.080
SCS-S	1.931	2.387	2.908	1.852	0.480	6.870	6.390
SCS-V	2.399	4.523	8.883	9.989	0.470	35.910	35.440
SPERRY1100-V	3.369	5.510	10.829	13.392	1.080	55.380	54.300

Table 4. Statistics of Reconstructed Data using Principal Components

* Ratio to Corresponding Entries in Table 3

System	1 Component			2 Components		
	Geometric*	Range*	Spearman	Geometric*	Range*	Spearman
ALLIANT-S-32	1.1034	0.1435	-0.6374	0.9896	0.9371	0.7529
ALLIANT-V-32	1.1887	0.0875	0.8802	0.9947	1.0961	0.9092
ALLIANT-S-64	1.1029	0.2658	-0.6334	0.9923	0.9229	0.7550
ALLIANT-V-64-P	1.2353	0.1092	0.8324	0.9980	0.7947	0.8424
AMDAHL5890-S	0.9417	1.5465	0.7193	0.9889	0.9146	0.6717
AMDAH1500VP-V	1.0589	0.7491	0.9070	0.9951	1.2184	0.9069
AMDAHL1200VP-V	1.1044	0.3714	0.9078	0.9986	0.7973	0.9085
AMDAHL1400VP-S	0.9595	1.2506	0.7667	0.9859	0.9567	0.7467
AMDAHL1400VP-V	1.1237	0.2350	0.9091	0.9987	0.5721	0.9114
APOLLO300-32	1.3136	0.8723	-0.0836	0.9924	0.1326	0.8375
APOLLO660-32	1.1889	1.6190	-0.2923	0.9923	0.5369	0.8023
APOLLO300-64	1.3581	2.9412	0.1445	1.0029	0.3915	0.7275
APOLLO660-64	1.2091	3.2907	-0.1481	0.9988	0.7308	0.7188
SUN3-64	1.1524	0.4620	-0.5684	0.9936	0.5907	0.6737
RIDGE32	1.1329	2.0640	-0.2762	0.9992	0.6884	0.6050
CDC875	1.0022	0.7627	0.8391	0.9960	0.8212	0.8439
CYBER176	1.0191	0.6126	0.8060	0.9922	0.8390	0.8220
CELERITY-32	1.1528	0.5157	-0.4560	0.9898	0.4688	0.6798
CONVEX-S-32	1.0753	0.1408	0.7620	1.0000	0.6539	0.7574
CONVEX-V-32	1.2359	0.1854	0.9210	0.9982	1.2539	0.9308
CONVEX-S-64	1.0891	0.0701	0.7192	1.0036	0.7051	0.6980
CONVEX-V-64	1.2221	0.1644	0.9184	0.9999	1.2243	0.9298
CRAY1-S	0.9815	0.8496	0.6901	0.9931	0.7509	0.6873
CRAY1-V	1.1348	0.5599	0.9765	1.0044	1.4540	0.9779
CRAYXMP-S	0.9716	0.9195	0.6951	0.9902	0.7598	0.6878
CRAYXMP-V	1.1517	0.5932	0.9792	1.0066	1.6600	0.9807
CRAYXMP-CFT-S	0.9860	0.8025	0.7272	0.9955	0.7312	0.7227
CRAYXMP-CFT-V	1.1438	0.5667	0.9835	1.0096	1.4742	0.9846
CRAY2-S	1.0006	0.7324	0.6558	0.9880	0.8386	0.6657
CRAY2-V	1.2059	0.3629	0.9740	1.0056	1.4706	0.9751
MICROVAX2	1.1742	1.5621	-0.6040	0.9998	0.8195	0.6669
VAX8800	1.0335	0.0154	-0.3762	0.9957	0.4041	0.4184
VAX8800-32	1.0324	0.2395	0.6077	0.9916	0.6810	0.6396
ELXSI6420	1.0575	0.1266	0.6546	0.9960	0.6591	0.6715
ETA205-S	1.0427	0.4808	0.7702	1.0045	0.6998	0.7684
ETA205-V	1.1362	0.1411	0.8879	1.0112	0.3658	0.8811
IBM3033	1.0366	0.3788	0.6525	0.9953	0.8853	0.6625
IBM3081	0.9852	0.9616	0.6930	0.9933	0.8357	0.6880
IBM3090-S	0.9422	1.7432	0.7305	0.9951	0.9426	0.6903
IBM3090-V	1.0292	0.6280	0.9125	0.9980	0.8165	0.9114
NECSX2-S	0.9199	1.2557	0.6352	0.9897	0.6312	0.5942
NECSX2-V	1.1278	0.3694	0.9670	1.0125	0.8266	0.9636
FPS264-64	1.0194	0.6801	0.7862	1.0088	0.7521	0.7845
HONEYWELDPDS-90	0.9958	0.6323	0.6095	0.9924	0.6564	0.6081
NASXL-60	0.9609	1.4039	0.7901	0.9853	1.1245	0.7799
SCS-S	1.0613	0.4022	0.7174	1.0004	0.8363	0.7275
SCS-V	1.2079	0.3204	0.9489	1.0097	1.4317	0.9455
SPERRY1100-V	1.1425	0.2680	0.7738	1.0001	0.8204	0.7917

Table 4. Statistics of Reconstructed Data using Principal Components (continued)

* Ratio to Corresponding Entries in Table 3

System	3 Components			4 Components		
	Geometric*	Range*	Spearman	Geometric*	Range*	Spearman
ALLIANT-S-32	0.9935	1.0300	0.9106	0.9923	1.0414	0.9087
ALLIANT-V-32	0.9944	1.0957	0.9096	0.9974	1.1746	0.9022
ALLIANT-S-64	0.9958	1.0337	0.9352	0.9953	1.0255	0.9360
ALLIANT-V-64-P	0.9920	0.8441	0.8994	0.9900	0.8250	0.9114
AMDAHL5890-S	0.9931	1.1172	0.9189	0.9975	1.0981	0.9238
AMDAH1500VP-V	0.9904	1.2343	0.9463	0.9982	1.4390	0.9723
AMDAHL1200VP-V	0.9915	0.8647	0.9669	1.0004	0.9677	0.9880
AMDAHL1400VP-S	0.9903	1.0618	0.9264	0.9936	1.0701	0.9179
AMDAHL1400VP-V	0.9907	0.6343	0.9680	0.9998	0.7039	0.9832
APOLLO300-32	0.9946	0.1437	0.8435	1.0007	0.3163	0.8281
APOLLO660-32	0.9946	0.6627	0.8713	0.9943	0.6584	0.8734
APOLLO300-64	1.0018	0.4921	0.6831	1.0056	1.0879	0.6859
APOLLO660-64	1.0005	0.9149	0.8312	1.0012	0.9513	0.8388
SUN3-64	0.9963	0.6297	0.7673	0.9975	0.6561	0.7750
RIDGE32	0.9999	0.7447	0.6190	1.0023	1.1488	0.6881
CDC875	0.9997	0.9204	0.9605	0.9993	0.9245	0.9619
CYBER176	0.9969	0.9537	0.9735	0.9978	0.9494	0.9720
CELERITY-32	0.9919	0.4979	0.7488	0.9903	0.4710	0.7493
CONVEX-S-32	1.0044	0.7609	0.9226	1.0044	0.7607	0.9224
CONVEX-V-32	1.0012	1.2467	0.9102	1.0011	1.2442	0.9102
CONVEX-S-64	1.0081	0.8323	0.9067	1.0082	0.8299	0.9056
CONVEX-V-64	1.0027	1.2200	0.9010	1.0023	1.2102	0.9002
CRAY1-S	0.9979	0.8575	0.9484	0.9992	0.8502	0.9407
CRAY1-V	1.0036	1.4557	0.9781	0.9989	1.3370	0.9867
CRAYXMP-S	0.9953	0.8601	0.9391	0.9962	0.8573	0.9281
CRAYXMP-V	1.0053	1.6646	0.9847	0.9995	1.4955	0.9904
CRAYXMP-CFT-S	1.0010	0.8102	0.9340	0.9994	0.8115	0.9416
CRAYXMP-CFT-V	1.0086	1.4769	0.9838	1.0017	1.2994	0.9914
CRAY2-S	0.9932	0.9474	0.8887	0.9931	0.9474	0.8887
CRAY2-V	1.0041	1.4758	0.9774	0.9978	1.3099	0.9889
MICROVAX2	1.0020	0.9340	0.8290	1.0029	0.9549	0.8280
VAX8800	0.9985	0.6632	0.7283	1.0027	0.8151	0.8363
VAX8800-32	0.9951	0.9279	0.8933	0.9986	0.9540	0.9451
ELXSI6420	0.9990	0.7470	0.7966	1.0019	0.8176	0.8352
ETA205-S	1.0092	0.7367	0.9114	1.0078	0.7303	0.9071
ETA205-V	1.0101	0.3663	0.8850	1.0048	0.3336	0.8889
IBM3033	0.9987	1.0993	0.8835	1.0021	1.1648	0.9016
IBM3081	0.9957	1.0609	0.9128	0.9990	1.0952	0.9324
IBM3090-S	0.9986	1.1408	0.8962	1.0021	1.0882	0.9306
IBM3090-V	0.9993	0.8198	0.9283	1.0022	0.8672	0.9509
NECSX2-S	0.9949	0.7598	0.9058	0.9971	0.7350	0.9142
NECSX2-V	1.0088	0.8346	0.9751	1.0041	0.7641	0.9790
FPS264-64	1.0120	0.7751	0.8141	1.0127	0.7810	0.8292
HONEYWELDPDS-90	0.9921	0.6521	0.6060	1.0026	1.2820	0.8554
NASXL-60	0.9917	1.2294	0.9647	0.9968	1.2997	0.9813
SCS-S	1.0057	0.9360	0.9292	1.0031	0.9443	0.9330
SCS-V	1.0102	1.4304	0.9436	1.0019	1.2280	0.9717
SPERRY1100-V	0.9961	0.8321	0.8415	1.0006	0.9076	0.8548

Table 5. Component Scores

System	Score 1	Score 2	Score 3
ALLIANT-S-32	0.7821	0.42106	1.98566
ALLIANT-V-32	1.4168	0.15247	0.94384
ALLIANT-S-64	0.6733	0.47188	1.85831
ALLIANT-V-64-P	3.0050	0.06459	0.34938
AMDAHL5890-S	7.5879	0.49322	1.91722
AMDAHL1500VP-V	22.6980	0.08722	0.39978
AMDAHL1200VP-V	34.7384	0.04568	0.25894
AMDAHL1400VP-S	8.1875	0.38367	2.01681
AMDAHL1400VP-V	39.1866	0.03546	0.22349
APOLLO300-32	0.0147	0.91159	1.69671
APOLLO660-32	0.1112	0.68069	1.60787
APOLLO300-64	0.0072	1.07103	0.99907
APOLLO660-64	0.0739	0.77801	1.47658
SUN3-64	0.3440	0.47283	1.63853
RIDGE32	0.2056	0.78035	1.19632
CDC875	4.0243	0.42162	1.78946
CYBER176	3.5725	0.37246	2.14754
CELERITY-32	0.2729	0.52357	1.50559
CONVEX-S-32	1.4067	0.41459	2.10933
CONVEX-V-32	3.5574	0.06571	1.63490
CONVEX-S-64	1.1665	0.42512	2.14767
CONVEX-V-64	2.4531	0.09266	1.58552
CRAY1-S	6.1201	0.37491	2.13668
CRAY1-V	16.4429	0.06193	0.80529
CRAYXMP-S	7.3995	0.35960	2.20850
CRAYXMP-V	24.3583	0.04372	0.71873
CRAYXMP-CFT-S	7.9166	0.31946	2.37490
CRAYXMP-CFT-V	24.2260	0.04826	0.77259
CRAY2-S	4.8969	0.34352	2.28666
CRAY2-V	16.4532	0.03538	0.71848
MICROVAX2	0.1810	0.62098	1.56825
VAX8800	0.9779	0.70629	1.63581
VAX8800-32	1.4197	0.56096	1.83011
ELXSI6420	1.2668	0.49606	1.67830
ETA205-S	5.0314	0.26380	2.13186
ETA205-V	9.7528	0.08137	0.78092
IBM3033	1.6318	0.53444	1.77900
IBM3081	2.5638	0.63028	1.47998
IBM3090-S	7.0721	0.53424	1.70343
IBM3090-V	11.0311	0.17673	1.13940
NECSX2-S	14.2086	0.39965	2.23597
NECSX2-V	62.8027	0.02995	0.46890
FPS264-64	6.9922	0.27576	1.61119
HONEYWELDPS-90	4.8738	0.36721	0.89663
NASXL-60	12.9588	0.29169	2.73871
SCS-S	2.7112	0.31604	2.39840
SCS-V	6.1840	0.06379	1.03702
SPERRY1100-V	7.2000	0.08270	0.47504

Table 6. Vectorization Statistics of Loops

Cluster	Loop	Q1	Q3	Vector	Cluster	Loop	Q1	Q3	Vector
1	102	.7225	4.935	4	2	101	1.7700	19.603	6
	104	.5435	3.653	4		201	1.7700	25.960	6
	204	.9388	6.393	4		301	1.7700	31.975	6
	105	.7518	5.598	0		203	1.3875	15.745	6
	205	.7353	5.818	0		303	1.4575	22.930	6
	305	.7600	5.863	0		304	1.1600	10.450	4
	106	.6190	3.253	3		107	2.1400	22.760	6
	111	.4820	4.473	0		207	2.2175	32.420	6
	211	.4563	5.045	0		307	2.1650	34.625	6
	311	.5075	5.088	0		108	1.3720	18.793	6
	116	.5260	3.555	0		208	1.4208	24.495	6
	216	.5198	3.515	0		308	1.3720	24.488	6
	316	.5203	3.515	0		109	1.8850	19.250	6
	117	.9308	8.325	0		209	2.0500	26.618	6
	217	.9308	7.905	0		309	2.0500	26.615	6
	317	.9270	7.905	0		110	.9210	9.550	4
	119	.8205	6.933	0		210	.9985	11.728	4
	219	.8330	7.253	0		310	.9885	11.728	4
	319	.8213	7.253	0		112	.7168	7.308	6
	120	1.2305	9.688	0		212	.7620	10.215	6
	220	1.2373	9.703	0		312	.7770	12.138	6
	320	1.2005	9.710	0		118	1.4700	13.925	6
	123	1.1593	10.735	2p		218	1.4725	22.690	6
	223	1.1908	10.923	2p		318	1.4725	22.128	6
	323	1.1600	10.915	2p		121	.9618	11.055	6
4	202	1.3250	8.028	4	3	221	.9265	11.120	6
	302	1.3075	8.028	4		321	.9618	11.088	6
	103	1.1918	9.430	6		122	.8180	7.078	6
	206	.9988	6.448	3		222	.8408	7.890	6
	306	1.0238	6.663	3		322	.8418	7.890	6
	113	.2910	3.103	4p		124	.3590	3.025	0
	213	.2910	3.430	4p		224	.3603	3.130	0
	313	.2910	3.488	4p		324	.3603	3.330	0
	114	.5425	5.163	4p					
	214	.6515	5.163	4p					
	314	.6385	5.163	4p					
	115	.7303	3.475	2					
	215	.8393	4.083	2					
	315	.8393	4.083	2					

Table 7. Cluster Component Scores

System	2 Clusters		4 Clusters			
	Score 1	Score 2	Score 1	Score 2	Score 3	Score 4
ALLIANT-S-32	0.620	0.880	0.721	0.908	0.339	0.515
ALLIANT-V-32	0.690	2.325	0.643	2.343	0.799	0.813
ALLIANT-S-64	0.553	0.741	0.635	0.760	0.347	0.462
ALLIANT-V-64-P	0.954	7.056	0.642	7.112	1.284	2.059
AMDAHL5890-S	6.058	8.532	7.693	8.761	3.783	4.234
AMDAH1500VP-V	7.600	51.584	7.014	50.774	26.168	7.980
AMDAHL1200VP-V	8.683	98.557	7.487	96.574	46.282	9.819
AMDAHL1400VP-S	5.948	9.865	7.487	10.232	3.300	4.291
AMDAHL1400VP-V	8.853	119.766	7.503	116.614	52.966	10.341
APOLLO300-32	0.016	0.013	0.019	0.013	0.039	0.011
APOLLO660-32	0.106	0.112	0.117	0.113	0.103	0.088
APOLLO300-64	0.008	0.007	0.008	0.007	0.033	0.006
APOLLO660-64	0.073	0.072	0.081	0.072	0.087	0.060
SUN3-64	0.287	0.372	0.341	0.379	0.210	0.222
RIDGE32	0.198	0.207	0.209	0.205	0.284	0.172
CDC875	3.109	4.521	3.439	4.715	1.260	2.956
CYBER176	2.715	4.019	3.122	4.193	1.130	2.399
CELERITY-32	0.234	0.295	0.263	0.301	0.159	0.201
CONVEX-S-32	1.097	1.565	1.321	1.637	0.400	0.909
CONVEX-V-32	1.237	7.196	1.548	7.916	0.408	0.972
CONVEX-S-64	0.912	1.292	1.136	1.351	0.338	0.711
CONVEX-V-64	0.973	4.542	1.204	4.951	0.340	0.772
CRAY1-S	4.679	6.848	5.459	7.103	2.290	3.935
CRAY1-V	4.908	38.974	5.128	42.919	2.140	5.114
CRAYXMP-S	5.622	8.294	6.530	8.619	2.620	4.800
CRAYXMP-V	6.120	65.846	6.317	73.316	2.590	6.546
CRAYXMP-CFT-S	5.607	9.312	6.610	9.879	1.580	5.007
CRAYXMP-CFT-V	6.549	60.457	6.472	68.266	1.553	8.224
CRAY2-S	3.700	5.511	4.195	5.738	1.640	3.322
CRAY2-V	3.857	46.622	3.900	52.073	1.673	4.263
MICROVAX2	0.159	0.194	0.190	0.196	0.135	0.118
VAX8800	0.934	0.967	1.121	0.973	0.839	0.681
VAX8800-32	1.249	1.478	1.509	1.506	0.888	0.932
ELXSI6420	1.083	1.318	1.174	1.345	0.713	0.997
ETA205-S	3.301	6.176	3.501	6.590	0.880	3.591
ETA205-V	3.620	18.763	3.971	20.793	0.857	3.772
IBM3033	1.332	1.816	1.675	1.853	0.983	0.926
IBM3081	2.218	2.770	2.613	2.802	1.923	1.692
IBM3090-S	5.749	7.841	7.092	8.067	3.340	4.270
IBM3090-V	5.757	16.690	6.606	17.596	3.321	4.881
NECSX2-S	11.205	15.528	13.435	16.172	4.590	9.204
NECSX2-V	13.989	179.995	12.875	203.412	4.540	19.069
FPS264-64	4.587	8.588	4.925	9.161	1.237	4.873
HONEYWELDPS-90	3.427	6.105	3.523	6.025	7.690	2.942
NASXL-60	8.831	15.750	12.121	16.520	3.706	5.685
SCS-S	1.926	3.171	2.276	3.377	0.480	1.744
SCS-V	1.996	13.346	2.088	14.910	0.477	2.261
SPERRY1100-V	2.445	16.138	2.083	16.633	4.651	3.042

Table 8. Statistics of Reconstructed Data using Cluster Components

* Ratio to Corresponding Entries in Table 3

System	2 Clusters			3 Clusters		
	Geometric*	Range*	Spearman	Geometric*	Range*	Spearman
ALLIANT-S-32	1.0000	0.4861	0.6752	1.0000	0.7485	0.8687
ALLIANT-V-32	1.0000	0.9231	0.9052	1.0000	0.8570	0.9050
ALLIANT-S-64	1.0000	0.4734	0.6546	1.0000	0.7109	0.8724
ALLIANT-V-64-P	1.0000	0.9153	0.8527	1.0000	0.7902	0.8669
AMDAHL5890-S	1.0000	0.6435	0.7080	1.0000	0.6813	0.7393
AMDAH1500VP-V	1.0000	1.5581	0.9199	1.0000	1.1525	0.9601
AMDAHL1200VP-V	1.0000	1.1162	0.9136	1.0000	0.7810	0.9828
AMDAHL1400VP-S	1.0000	0.6852	0.7371	1.0000	0.7596	0.7755
AMDAHL1400VP-V	1.0000	0.8093	0.9165	1.0000	0.5489	0.9786
APOLLO300-32	1.0000	0.2127	0.4794	1.0000	0.2284	0.4674
APOLLO660-32	1.0000	0.5154	0.6455	1.0000	0.4045	0.8575
APOLLO300-64	1.0000	0.6024	0.6232	1.0000	1.2886	0.3474
APOLLO660-64	1.0000	0.9654	0.4955	1.0000	0.7052	0.6672
SUN3-64	1.0000	0.2865	0.4955	1.0000	0.3696	0.7272
RIDGE32	1.0000	0.7042	0.4617	1.0000	0.7990	0.4410
CDC875	1.0000	0.4825	0.8189	1.0000	0.7456	0.9424
CYBER176	1.0000	0.4368	0.7939	1.0000	0.6704	0.9409
CELERITY-32	1.0000	0.2446	0.5555	1.0000	0.3551	0.7770
CONVEX-S-32	1.0000	0.3288	0.7310	1.0000	0.6405	0.9089
CONVEX-V-32	1.0000	0.9670	0.9220	1.0000	1.1683	0.9321
CONVEX-S-64	1.0000	0.3612	0.6620	1.0000	0.7020	0.8861
CONVEX-V-64	1.0000	0.9659	0.9228	1.0000	1.1710	0.9309
CRAY1-S	1.0000	0.4177	0.6802	1.0000	0.5443	0.7851
CRAY1-V	1.0000	1.5999	0.9793	1.0000	1.7540	0.9793
CRAYXMP-S	1.0000	0.4137	0.6695	1.0000	0.5369	0.7648
CRAYXMP-V	1.0000	1.9515	0.9824	1.0000	2.1305	0.9828
CRAYXMP-CFT-S	1.0000	0.4375	0.7108	1.0000	0.6473	0.8907
CRAYXMP-CFT-V	1.0000	1.5529	0.9841	1.0000	1.8692	0.9827
CRAY2-S	1.0000	0.4260	0.6252	1.0000	0.6006	0.7611
CRAY2-V	1.0000	1.6586	0.9784	1.0000	1.8297	0.9789
MICROVAX2	1.0000	0.5768	0.4119	1.0000	0.6573	0.7236
VAX8800	1.0000	0.0590	0.3337	1.0000	0.1485	0.5744
VAX8800-32	1.0000	0.2595	0.6095	1.0000	0.4616	0.7744
ELXSI6420	1.0000	0.2296	0.6527	1.0000	0.4257	0.7688
ETA205-S	1.0000	0.4195	0.7714	1.0000	0.6215	0.8987
ETA205-V	1.0000	0.3207	0.8904	1.0000	0.4000	0.8816
IBM3033	1.0000	0.5379	0.6446	1.0000	0.6778	0.6895
IBM3081	1.0000	0.5782	0.6639	1.0000	0.5589	0.6544
IBM3090-S	1.0000	0.7075	0.7230	1.0000	0.8025	0.7855
IBM3090-V	1.0000	0.6630	0.9103	1.0000	0.6941	0.9127
NECSX2-S	1.0000	0.3937	0.6209	1.0000	0.4909	0.7525
NECSX2-V	1.0000	0.9208	0.9665	1.0000	1.0223	0.9656
FPS264-64	1.0000	0.4849	0.7994	1.0000	0.6945	0.7940
HONEYWELDPS-90	1.0000	0.5407	0.6167	1.0000	0.7861	0.5755
NASXL-60	1.0000	0.6719	0.7561	1.0000	0.7956	0.8414
SCS-S	1.0000	0.4535	0.7178	1.0000	0.8016	0.9018
SCS-V	1.0000	1.3033	0.9517	1.0000	1.6378	0.9558
SPERRY1100-V	1.0000	1.0163	0.8043	1.0000	0.8664	0.8490

Table 8. Statistics of Reconstructed Data using Cluster Components (continued)

* Ratio to Corresponding Entries in Table 3

System	4 Clusters			5 Clusters		
	Geometric*	Range*	Spearman	Geometric*	Range*	Spearman
ALLIANT-S-32	1.0000	0.8504	0.8772	1.0000	0.8855	0.8751
ALLIANT-V-32	1.0000	0.7910	0.9029	1.0000	0.6632	0.8953
ALLIANT-S-64	1.0000	0.8545	0.9002	1.0000	0.8553	0.9002
ALLIANT-V-64-P	1.0000	0.6544	0.9173	1.0000	0.8431	0.9049
AMDAHL5890-S	1.0000	1.1312	0.9345	1.0000	1.0712	0.9523
AMDAH1500VP-V	1.0000	1.2484	0.9754	1.0000	1.0842	0.9741
AMDAHL1200VP-V	1.0000	0.8112	0.9893	1.0000	0.8774	0.9899
AMDAHL1400VP-S	1.0000	1.0917	0.9224	1.0000	0.9645	0.9454
AMDAHL1400VP-V	1.0000	0.5627	0.9864	1.0000	0.7230	0.9841
APOLLO300-32	1.0000	0.2462	0.7460	1.0000	0.2453	0.7353
APOLLO660-32	1.0000	0.5298	0.8440	1.0000	0.5312	0.8340
APOLLO300-64	1.0000	1.2736	0.7258	1.0000	1.2288	0.7533
APOLLO660-64	1.0000	0.8522	0.7989	1.0000	0.8521	0.7989
SUN3-64	1.0000	0.5644	0.7891	1.0000	0.8403	0.8502
RIDGE32	1.0000	1.0004	0.6564	1.0000	1.0365	0.6747
CDC875	1.0000	0.7546	0.9443	1.0000	0.8406	0.9486
CYBER176	1.0000	0.7218	0.9632	1.0000	0.7561	0.9628
CELERITY-32	1.0000	0.4190	0.7696	1.0000	0.7402	0.8257
CONVEX-S-32	1.0000	0.7108	0.9279	1.0000	0.7169	0.9315
CONVEX-V-32	1.0000	1.3530	0.9096	1.0000	1.3539	0.9096
CONVEX-S-64	1.0000	0.8287	0.9238	1.0000	0.8301	0.9225
CONVEX-V-64	1.0000	1.3428	0.9021	1.0000	1.2456	0.9066
CRAY1-S	1.0000	0.6328	0.9143	1.0000	0.6063	0.9167
CRAY1-V	1.0000	1.7415	0.9793	1.0000	1.1914	0.9881
CRAYXMP-S	1.0000	0.6145	0.8826	1.0000	0.6297	0.8821
CRAYXMP-V	1.0000	2.0948	0.9839	1.0000	1.4904	0.9863
CRAYXMP-CFT-S	1.0000	0.6803	0.9051	1.0000	0.6603	0.9282
CRAYXMP-CFT-V	1.0000	1.5839	0.9896	1.0000	1.2167	0.9876
CRAY2-S	1.0000	0.6468	0.8391	1.0000	0.6287	0.8430
CRAY2-V	1.0000	1.7499	0.9818	1.0000	1.1928	0.9861
MICROVAX2	1.0000	1.0851	0.8562	1.0000	1.0885	0.8556
VAX8800	1.0000	0.8057	0.8118	1.0000	0.8343	0.8587
VAX8800-32	1.0000	0.8177	0.9151	1.0000	0.8877	0.9290
ELXSI6420	1.0000	0.4576	0.7958	1.0000	0.6107	0.8295
ETA205-S	1.0000	0.5557	0.7996	1.0000	0.5110	0.7905
ETA205-V	1.0000	0.3773	0.8926	1.0000	1.1272	0.9198
IBM3033	1.0000	1.2658	0.9166	1.0000	1.2209	0.9165
IBM3081	1.0000	1.1812	0.9513	1.0000	1.1186	0.9590
IBM3090-S	1.0000	1.1744	0.9303	1.0000	1.1957	0.9331
IBM3090-V	1.0000	0.8143	0.9509	1.0000	1.0878	0.9590
NECSX2-S	1.0000	0.5693	0.8942	1.0000	0.6292	0.9057
NECSX2-V	1.0000	0.8513	0.9734	1.0000	1.0966	0.9761
FPS264-64	1.0000	0.6367	0.7650	1.0000	0.7279	0.7851
HONEYWELDPS-90	1.0000	0.7983	0.7892	1.0000	0.9471	0.9001
NASXL-60	1.0000	1.1732	0.9793	1.0000	1.0808	0.9761
SCS-S	1.0000	0.8092	0.9062	1.0000	0.8051	0.9198
SCS-V	1.0000	1.4492	0.9625	1.0000	1.3654	0.9653
SPERRY1100-V	1.0000	0.8201	0.8566	1.0000	0.6755	0.8466





Figure 1a. Residual Plot for Loop Length 1






Figure 1c. Residual Plot for Loop Length 3




Figure 2. Plot of Summary Scores

U.S. DEPT. OF COMM. BIBLIOGRAPHIC DATA SHEET <i>(See instructions)</i>	1. PUBLICATION OR REPORT NO. NISTIR 88-4019	2. Performing Organ. Report No.	3. Publication Date January 1989
4. TITLE AND SUBTITLE On the Analysis of Computer Performance Data			
5. AUTHOR(S) Jack C. M. Wang, John M. Gary, and Hari K. Iyer			
6. PERFORMING ORGANIZATION (If joint or other than NBS, see instructions) NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY NATIONAL BUREAU OF STANDARDS U.S. DEPARTMENT OF COMMERCE GAITHERSBURG, MD 20899			7. Contract/Grant No. 8. Type of Report & Period Covered
9. SPONSORING ORGANIZATION NAME AND COMPLETE ADDRESS (Street, City, State, ZIP)			
10. SUPPLEMENTARY NOTES <input type="checkbox"/> Document describes a computer program; SF-185, FIPS Software Summary, is attached.			
11. ABSTRACT (A 200-word or less factual summary of most significant information. If document includes a significant bibliography or literature survey, mention it here) This paper is devoted to an analysis of the data from the Livermore loops benchmark. We will show that in a general predictive sense the dimension of this data is rather small; perhaps between two and five. Two techniques are used to reduce the 72 loop timings for each machine to a few scores which characterize the machine. The first is based on a principal component analysis, the second on a cluster analysis of the loops. The validity of the reduction of the data to a lesser dimension is checked by various methods.			
12. KEY WORDS (Six to twelve entries; alphabetical order; capitalize only proper names; and separate key words by semicolons) benchmarks; computers; performance; Livermore loops; principal components; clusters			
13. AVAILABILITY <input checked="" type="checkbox"/> Unlimited <input type="checkbox"/> For Official Distribution. Do Not Release to NTIS <input type="checkbox"/> Order From Superintendent of Documents, U.S. Government Printing Office, Washington, D.C. 20402. <input checked="" type="checkbox"/> Order From National Technical Information Service (NTIS), Springfield, VA. 22161			14. NO. OF PRINTED PAGES 44 15. Price \$12.95

